

042390.P13945

Docket No.: 042390.P13945  
Express Mail No.: EL802885258US

UNITED STATES PATENT APPLICATION

FOR

DYNAMICALLY RECONFIGURING CLOCK DOMAINS ON A CHIP

INVENTOR:

Sunil C. Chaudhari

PREPARED BY:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP  
12400 Wilshire Blvd., 7th Floor  
Los Angeles, CA 90025-1026  
(714) 557-3800

## DYNAMICALLY RECONFIGURING CLOCK DOMAINS ON A CHIP

### BACKGROUND

#### 1. FIELD

Embodiments of the invention relate to the field of clock architectures. More particularly, embodiments of the invention relate to dynamically reconfiguring clock domains on a chip.

#### 2. DESCRIPTION OF RELATED ART

Since the advent of the integrated circuit (IC), circuit components have become smaller and smaller. An IC may include millions of components packed into an incredibly small package. With each new generation of smaller integration, more functionality, and therefore more value, can be derived from ICs. Reliably manufacturing these highly integrated ICs, however, presents significant design challenges.

In particular, designing ICs that meet timing constraints can be particularly difficult. An IC may include tens of thousands of registers that need to be connected to one or more clock sources. For each clock "tick", or clock transition, thousands of registers have to operate in concert. A complex network is needed to propagate the clock signal to each of the registers. If the difference in propagation delay through two different paths in the network is too large or too small, errors may occur that can cause the entire IC to fail. Those skilled in the art will be familiar with numerous processes for synthesizing clock networks, or clock tree solutions.

Conventional integrated circuits (ICs) use a clock signal and branch it out through a series of buffers to form a plurality of clock signals. The structure of the branching of the clock signal is called a "clock tree." One or more clock trees can be present in a single IC. The clock signals at any level of the clock tree are sent to

various synchronous components of the IC to coordinate the functions of these components. For various reasons, however, any two clock signals, even at the same level of the same clock tree, may be slightly different or offset from each other. This difference in clock signals is called "clock skew".

- 5           Clock skew has several causes. For example, the buffers between levels in the clock tree typically introduce a delay between their input and output clock signals, so clock signals at different levels of the clock tree are usually naturally skewed from each other. Additionally, the load experienced by one clock signal may introduce a delay into the clock signal different from that of another load on another clock signal.
- 10          Furthermore, changes in temperature, different applied voltages and differing semiconductor fabrication processes can affect the clock skew.

- Common IC fabrication techniques try to minimize clock skew by resizing buffers in the clock tree to move some of the clock signals forward or backward or by adding redundant loads to the circuits to balance the loading of the clock tree. Either
- 15          technique alters the delay of some of the clock signals by a specified amount that is determined by analysis of the timing of the clock signals. With temperature changes, different applied voltages, differing silicon fabrication processes and/or inadequate tolerances in the silicon fabrication processes, the clock skew can even vary significantly from the nominal value. Due to this variation in clock skew, such ICs will
- 20          often fail speed testing and not meet performance targets. As a result, IC fabrication yields will be low and costs will be high.

- Further, as ICs continue to become more complex, having tens of thousand of registers which may be clocked by several different source clocks, at several different clock frequencies, through gated clocks, inverted clocks, etc., the processing time and
- 25          expense required to meet continually more stringent design constraints using known approaches is becoming increasingly prohibitive. Moreover, due to the ever increasing complexity of newer ICs, low skew clock distribution becomes more and more difficult

to achieve. Normally, when an IC is designed, clock domains are chosen with a size that can achieve the desired clock skew, and synchronization logic is added on the signals that cross the clock domains. This synchronization logic adds additional latency on the signals that cross the clock domain partitions.

- 5            This clock domain partitioning is typically done before fabricating the chip, since choosing a large clock domain could result in clock skew problems in the real silicon that is very expensive to fix with another revision of the IC. The difficulty is that the timing models that are used to design the clock tree are usually not accurate enough with respect to the real silicon to design a large clock tree with little clock
- 10   skew.

## BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 illustrates a process to design dynamically reconfigurable clock domains on an integrated circuit, according to one embodiment of the invention.

Figure 2 shows a block diagram of an exemplary integrated circuit having  
5 multiple clock signal paths in order to implement dynamically reconfigurable clock domains, according to one embodiment of the invention.

Figure 3 shows a block diagram illustrating a more detailed view of a circuit component of Figure 2, utilized in implementing dynamically reconfigurable clock domains, according to one embodiment of the invention.

10 Figure 4 illustrates a process for providing dynamically reconfigurable clock domains on an integrated circuit, according to one embodiment of the invention.

Figure 5 shows a block diagram of a very basic computer system configuration that provides an exemplary environment in which an integrated circuit having dynamically reconfigurable clock domains, according to one embodiment of the  
15 invention, can be practiced.

## DESCRIPTION

In the following description, various embodiments of the invention will be described in detail. However, such details are included to facilitate understanding of the invention and to describe exemplary embodiments for employing the invention.

5 Such details should not be used to limit the invention to the particular embodiments described because other variations and embodiments are possible while staying within the scope of the invention. Furthermore, although numerous details are set forth in order to provide a thorough understanding of the embodiments of the invention, it will be apparent to one skilled in the art that these specific details are not required in order

10 to practice the embodiments of the invention. In other instances details such as, well-known methods, types of data, protocols, procedures, components, electrical structures and circuits, are not described in detail, or are shown in block diagram form, in order not to obscure the invention. Furthermore, embodiments of the invention will be described in particular embodiments but may be implemented in hardware, software,

15 firmware, middleware, or a combination thereof.

Embodiments of the invention relate generally to an integrated circuit that includes clock generation circuitry which generates a master clock signal and at least one other clock signal. The master clock signal and the other clock signal are transmitted through a clock distribution tree to a circuit component. In a default mode,

20 the circuit component receives the master clock signal at a first component block to create a first time domain for the first component block of the circuit component and the circuit component receives the other clock signal at a second component block to create a second time domain for the second component block of the circuit component.

Embodiments of the invention may further include bypass logic which creates a

25 bypass path to allow the second component block to receive the master clock signal such that the first and second components are controlled by a common clock domain in response to the master clock signal. Accordingly, the clock domain of the second

component block of the circuit component is the same as the clock domain of the first component block such that signals can be transferred between the clock domains with reduced latency- as opposed to the latency typically caused by the required synchronization logic of the default mode. Moreover, embodiments of the invention  
5 relate to processes to design dynamically reconfigurable clock domains (e.g. the default and bypass modes) on an integrated circuit, as previously described.

As shown in Figure 1, Figure 1 illustrates a process 100 to design dynamically reconfigurable clock domains on an integrated circuit, according to one embodiment of the invention. At block 110 of process 100, the logic of the plurality of circuit  
10 components on an integrated circuit are designed in such a way as to achieve a desired clock skew less than a predetermined minimum. The process of designing integrated circuits in view of obtaining desired clock skews is well known in the art. Particularly, as is known in the art, the design of the integrated circuit involves partitioning the chunks of logic of the integrated circuit in such a way so that the sum of the minimum  
15 clock to Q delay of the flip-flop and the minimum combinational delay path is greater than the sum of the maximum clock skew for any path and the maximum hold time of the flip-flop in the particular technology. This is to reduce the probability of having hold time problems while reducing the degradation in frequency due to the clock skew. The chunks of combinational logic typically contain many flip-flops (e.g. 10 thousand  
20 flip-flops, 20 thousand flip-flops, etc.).

Next, at block 120 of process 100, a clock distribution tree is designed in such a way as to achieve the desired clock skew (previously determined in block 110) across all combinable circuit components, component blocks of the circuit components, and combinational logic blocks of the circuit components for the integrated circuit (as will  
25 be discussed). Particularly, as will be discussed in more detail later, the clock distribution tree is designed with the goal of achieving the desired clock skew across all partitions that may possibly be combined together, later, by software configuration. Although, determining clock distribution trees for ICs is well know in the art,

embodiments of the invention determine the clock distribution tree for all partitions of the IC that may possibly be combined together later by software configuration.

Further, at block 130 of process 100, two paths for signals that cross clock domain partitions are defined (as will be discussed in detail later with reference to  
5 Figures 2 and 3), including: 1) a default path; and 2) a bypass path. In one embodiment, these paths can be selected via software. By default, the IC utilizes the default path, which contains synchronization logic that is used when the IC is running in default mode with separate clock domains. The bypass path includes  
10 synchronization bypass logic that is used when the chip is reconfigured to combine clock domains. The logic that produces and receives the signals across the clock domains may be designed to work in both handshaking modes for the synchronization logic of the default mode and the synchronization bypass logic of the bypass mode. Also, the setup and hold time analysis performed to create the clock distribution tree may be done for the bypass mode as well as for the default mode.

15 As will be discussed, when the IC powers up, all the clock domains are separate by default. In one embodiment, once it is proven that the achieved clock skew between the clock domains in the chip is less than the original goal, software can be used to write a bit to combine the clock domains which results in the bypass mode being selected (e.g. utilizing bypass paths). In the bypass mode, clock domains for a signal  
20 are combined into one, such that there is no latency on the signals that cross clock domains. Thus, signals that cross different component blocks of a circuit component of an IC are synchronous with one another thereby reducing latency.

As shown in Figure 2, Figure 2 illustrates a block diagram of an exemplary integrated circuit having multiple clock signal paths in order to implement dynamically  
25 reconfigurable clock domains, according to one embodiment of the invention. The exemplary integrated circuit 200 that implements dynamically reconfigurable clock domains is designed in accordance with the process 100, previously discussed. As



shown in Figure 2, integrated circuit 200 includes clock generation circuitry 204, a plurality of switches 210 such as multiplexers (e.g. multiplexer 210-2, multiplexer 210-3, etc., through multiplexer 210-N), and circuit components 206 (e.g. circuit component-1 206-1, circuit component-2 206-2, etc., through circuit component-N 206-N). It should be appreciated that as shown in Figure 2, the integrated circuit 200 is scalable to any number of components N.

The integrated circuit 200 includes clock generation circuitry 204 that generates a plurality of clock signals (CLK1-N) that are coupled to circuit components 206-1 through 206-N, respectively, to generate a clock distribution tree 205 of multiple clock signals paths. Specifically, clock generation circuitry 204 generates a Master clock signal (CLK1) and a plurality of other clock signals (also termed circuit component clock signals): CLK2, CLK3, CLK4, CLK5, and CLKN, as shown in Figure 2. In this embodiment, the plurality of clock signals (CLK1-N) generated by the clock generation circuitry 204 are each phase and frequency independent of one another. Each circuit component 206-1 through 206-N of the IC 200, respectively, may be any type of circuit component typically found in an integrated circuit that requires a clock signal. For example, each circuit component may be a bus, a processor, a memory, or any type of computer circuit component.

As shown in Figure 2, in this embodiment, the Master clock signal (CLK1) is directly coupled to circuit component-1 206-1 and is directly coupled to a first component block 207 of each of the circuit components 206-2 through 206-N, respectively. Thus, a first component block 207 of each circuit component 206-2 through 206-N is timed in accordance with the Master clock signal (CLK1).

As shown in Figure 2, each multiplexer 210-2 through 210-N has two inputs: the Master clock signal (CLK1) and its respective circuit component clock signal (e.g. CLK2 through CLKN). Multiplexers 210-2 through 210-N are by default set to 0 such that a second component block 209 of each respective circuit component 206 is timed

by its respective circuit component clock signal (e.g. CLK2 through CLKN). This is termed the default path. However, each multiplexer 210-2 through 210-N is responsive to a respective Master select signal (e.g. Master select signal-2 through Master select signal-N) which, when chosen, selects a bypass mode that instead applies the Master clock (CLK1) to the second component block 209 of the circuit component 206. In this way, a bypass path is created such that signals that traditionally cross time domains in the circuit component (e.g. from the first component block 207 to the second component block 209), resulting in increased latency, can instead operate at the same clock rate (e.g. the clock domains are combined) so that latency is reduced. Thus, signals that cross different component blocks of a circuit component of an IC are instead synchronous with one another thereby reducing latency.

Turning now to Figure 3, Figure 3 shows a block diagram 300 illustrating a more detailed view of circuit component-2 206-2 of Figure 2, utilized in implementing dynamically reconfigurable clock domains, according to one embodiment of the invention. Particularly, Figure 3 shows a circuit component that can operate in either default mode (via a default path 311) or in a bypass mode (via a bypass path 312). As should be appreciated, this description is applicable to any of the circuit components 206-2 through 206-N.

As shown in Figure 3, a signal crossing from a first time domain of the first component block 207 of the circuit component-2 206-2 to the second component block 209 of circuit component-2 having a second time domain, travels along a path that is timed using flip-flop 310, which utilizes Master clock (CLK1). For example, the Master clock (CLK1) signal may correspond to the clock signal of a system bus (e.g. operating at 200 MHz) and the circuit component-2 206-2 may be a specialized core processor (e.g. operating at 218 MHz by default (e.g. CLK2)). The signal crosses into the second component block 209 of circuit component-2 206-2 and crosses clock domains along the default path 311. Synchronization logic is then added on to the

signal that cross the clock domains into the second component block 209 of circuit component-2 206-2.

In this example, the synchronization logic includes a double clock synchronizer 320 including a first flip-flop 322 timed at CLK2 and a second flip-flop 324 also timed  
5 at CLK2. The signal is clocked by the first and second flip-flops 322 and 324 such that it is translated from the first time domain of the first component block 207 to the second time domain of the second component block 209 and is then passed on to combination logic 340 for processing. Unfortunately, this use of synchronization logic 320 may add additional latency on signals that cross the clock domain partitions.

10 According to one embodiment of the invention, instead of a signal having to cross time domains, the time domains can instead be combined, and the signal can follow a bypass path 312. In this instance, a signal crosses from a combined time domain from the first component block 207 of the circuit component-2 206-2 to the second component block 209 of circuit component-2 having the same combined time  
15 domain. For example, the combined time domains can be set to the Master clock signal (CLK1). Thus, the signal is first timed using flip-flop 310 by Master clock (CLK1). The signal then crosses into and through the second component block 209 of circuit component-2 206-2 utilizing the same combined clock domain (e.g. the same time domain as Master clock (CLK1)) along the bypass path 312.

20 In order to accomplish this, a bypass switch 330, e.g. a bypass multiplexer, has to be switched on such that the bypass path 312 is activated. Typically, the bypass multiplexer 330 is by default set to 0 such that the default path 311 is used. However, the bypass multiplexer 330 is activated in response to a bypass path activation signal 314 such that the bypass 312 is used instead of the default path 311. The bypass path  
25 activation signal 314 is sent when the corresponding master select signal of the corresponding multiplexer 210 is activated indicating that the bypass mode has been chosen for the given circuit component (e.g. in this example circuit component-2 with

the master select signal-2). The combination of the master select signal, switch 210, bypass path activation signal 314, and bypass switch 330 can generally be referred to as the bypass logic. Thus, when the bypass path 312 is chosen the Master clock (CLK1) is applied to signals crossing over into the combinational logic 340 of the second  
5 component block 209 of the circuit component 206 for processing. In one embodiment, the master select signal and bypass path activation signal can be selected using software, as will be discussed.

Therefore, the synchronization logic (e.g. the double clock synchronizer 320) that is typically added on to the signal when it crosses clock domains is totally avoided  
10 and the resultant latency to the signal is also avoided. In this way, a bypass path is created such that signals which traditionally cross time domains, resulting in increased latency, can instead operate at the same clock rate (e.g. the clock domains can be combined) so that latency is reduced. Thus, signals that cross different component blocks of a circuit component of an IC are instead synchronous with one another  
15 thereby reducing latency.

It should be appreciated that the integrated circuit (IC) 200, clock distribution tree 205, combinational logic 340, default paths, bypass paths, etc., of Figures 2 and 3, only illustrate one example of an IC that can be created utilizing the process 100, previously discussed. ICs with a multitude of different configurations can be created  
20 utilizing process 100. Further, it should be appreciated by those skilled in the art that to properly implement the IC 200 the clock distribution tree 205 needs to be designed in such a way as to achieve the desired clock skew across all combinable circuit components 206-1 through 206-N, combinable component blocks (e.g. 207 and 209), and combinational logic blocks 340 of the circuit components of the IC 200.  
25 Particularly, the clock distribution tree needs to be designed with the goal of achieving the desired clock skew across all partitions 207, 209 of all the circuit components 206 that may possibly be combined together later by software configuration in accordance with block 120 of process 100. Further, as previously shown in Figures 2 and 3,

process block 130 of process 100 requires that two paths for signals that cross clock domain partitions are defined including: 1) a default path 311; and 2) a bypass path 312. In one embodiment, these paths can be selected via software.

As previously discussed, by default, the IC 200 utilizes the default path 311,  
5 which contains synchronization logic 320 that is used when the IC is running in default mode with separate clock domains. However, as previously discussed, the bypass path 312 includes synchronization bypass logic that is used when the IC 200 is reconfigured to combine clock domains. As previously discussed with reference to process 100, the logic that produces and receives the signals across the clock domains may be designed  
10 to work in both handshaking modes for the synchronization logic 320 of the default mode and the synchronization bypass logic of the bypass mode. Also, the setup and hold time analysis performed to create the clock distribution tree 205 may be done for the bypass mode as well as for the default mode.

As will be discussed, when the IC powers up, all the clock domains are separate  
15 by default. In one embodiment, once it is proven that the achieved clock skew between the clock domains in the chip is less than the original goal, software can be used to write a bit to combine the clock domains which results in the bypass mode being selected (e.g. utilizing bypass paths). As illustrated previously, in the bypass mode, clock domains for a signal are combined into one, such that there is no latency on  
20 signals that cross the clock domains. Thus, signals that cross different component blocks of a circuit component of an IC are synchronous with one another thereby reducing latency.

Turning now to Figure 4, Figure 4 illustrates a process 400 for providing dynamically reconfigurable clock domains on an integrated circuit 200, according to  
25 one embodiment of the invention. At block 402 of process 400, the integrated circuit is powered up. All of the clock domains for all of the circuit components are set separately by default to run in default mode (e.g. to utilize their default paths) (block

404). For example, in the example of the integrated circuit 200 of Figures 2 and 3, all of the circuit components 206 are set to run in their default modes using their default paths 311 based on their respective clock domains (e.g. CLKs2-N).

Next, at block 406, it is determined whether the IC is functional in view of clock skew. Typically, this is done by the manufacturer of the IC utilizing test equipment. If the IC is found not to be functional in view of clock skew, due to a given circuit component or all of the circuit components, then the IC is designated as non-functional (block 408). On the other hand, if it is proven that the achieved clock skew between the clock domains in the chip is less than the original goal, and the IC is therefore functional, software can be used to write a bit to combine the clock domains which results in the bypass mode being selected (e.g. utilizing bypass paths) (block 409). Next, the process 400 verifies whether the chip is still functional with combined clock domains (block 410). If the IC is not functional with the combined clock domains, it is designated as only working in default mode (block 412). However, if the IC is still functional with combined clock domains it is then designated as being functional in bypass mode (block 414). A user of the IC (or the manufacturer) can choose whether to use the default mode or the bypass mode with combined clock domains (block 420).

It should be noted that combining clock domains can be done for all of the circuit components (e.g. all of the circuit components are selected) of the IC or, in some embodiments, only to selected circuit components of the IC. This is accomplished by selecting the bypass paths of the selected circuit components (block 410), as previously discussed in detail with reference to Figures 2 and 3.

Further, in one embodiment, the selection of bypass modes for all of the clock domains or only specific clock domains can be done by software. The software can be used to set a bit to one in a clocking register for each circuit component selected for bypass mode to designate that the clock domains are to be combined for that circuit

component 206 (e.g. the bypass path is chosen such that the selected circuit component 206 is timed with the Master clock CLK1). This results in the Master clock select signal being driven to the appropriate multiplexer 210 for the selected circuit component 206 and the bypass signal 314 being driven to the appropriate bypass  
5 multiplexer 330 in the selected circuit component 206. In this way, the different component blocks of the circuit component operate synchronously with the Master clock signal. Particularly, a signal clocked at a first rate does not have to go through synchronization logic 320 to be converted from one time domain to another time domain. Thus, signals that cross different component blocks of a circuit component of  
10 an IC are instead synchronous with one another thereby reducing latency.

Further, after testing proves that the IC is functional in view of the clock skew and is functional in default modes and bypass modes, a user of the IC (or the manufacturer) can choose whether to use the default mode or the bypass mode with combined clock domains (block 420). Typically, the IC is automatically set to operate  
15 in default mode (i.e. utilizing default paths). However, a user (or the manufacturer) can select that all of the circuit components of the IC or certain components of the IC operate with combined clock domains (e.g. utilizing bypass paths) by using software to select combined clock domains, as previously discussed. For example, even if the IC is functional in default mode, the user may wish to program the chip using software to  
20 combine the clock domains of all of the circuit components of the IC or only certain circuit components of the IC such that they run with less latency on the signals that communicate between the clock domains, thus resulting in higher performance. However, after circuit components of the IC are set to operate in bypass mode, they cannot be changed back to operate in default mode again, unless the whole IC is reset  
25 (e.g. a phase locked loop (PLL) reset is performed).

Referring briefly to Figure 5, Figure 5 shows a block diagram of a very basic computer system configuration 500 that provides an exemplary environment in which an integrated circuit 200 having dynamically reconfigurable clock domains, according

to one embodiment of the invention, can be practiced. A user can utilize the integrated circuit 200 with a vast multitude of different computer systems 500. A basic computer system configuration having a processor 502 and a memory 504 is shown. Particularly, the memory 504 may store software to write bits into a clocking register such that the  
5 system 500 under the control of processor 502 may utilize the software to provide the user (or the manufacturer) with the function of choosing to combine clock domains for all of the circuit components of the IC 200 or only selected circuit components of the IC 200, as previously discussed. Thus, using the system configuration 500 a user can utilize the IC 200 having dynamically reconfigurable clock domains. As previously  
10 discussed, when clock domains are combined for a circuit component, signals that cross different component blocks of the circuit component of the IC are synchronous with one another thereby reducing latency.

For the purposes of the present specification, the term “processor” or “CPU” refers to any machine that is capable of executing a sequence of instructions and shall  
15 be taken to include, but not be limited to, general purpose microprocessors, special purpose microprocessors, application specific integrated circuits (ASIC), multi-media controllers, signal processors and microcontrollers, etc. The memory 504 can include any memory device adapted to store digital information, such as static random access memory (SRAM), dynamic random access memory (DRAM), synchronous dynamic  
20 random access memory (SDRAM), and double data rate (DDR) SDRAM or DRAM, etc. Thus, in one embodiment, memory 504 include volatile memory. Further, memory 504 can also include non-volatile memory such as read-only memory (ROM). It should be appreciated that the exemplary illustration of the computer system configuration 500 is only an exemplary illustration of an environment in which  
25 embodiments of the invention can be practiced.

According to embodiments of the invention, a user or the manufacturer may decide whether or not to combine clock domains. For example, if a user chooses not to combine clock domains, then the IC 200 continues to operate using the default time



domain. However, if a user chooses to combine clock domains (e.g. for all of the circuit components of the IC or only certain circuit components of the IC), then software stored in memory 504 under the control of processor 502 may be utilized to write a bit in a clocking register for the selected circuit component in order to combine  
5 the respective clock domains for the circuit component, as previously discussed. This can be done for all of the circuit components or selected circuit components to select the bypass paths of the circuit components such that the clock domains for different component blocks of circuit components are combined, as previously discussed in detail with reference to Figures 2 and 3.

10 While embodiments of the present invention and its various functional components have been described in particular embodiments, it should be appreciated the embodiments of the present invention can be implemented in hardware, software, firmware, middleware or a combination thereof and utilized in systems, subsystems, components, or sub-components thereof. When implemented in software or firmware,  
15 the elements of the present invention are the instructions/code segments to perform the necessary tasks. The program or code segments can be stored in a machine readable medium (e.g. a processor readable medium or a computer program product), or transmitted by a computer data signal embodied in a carrier wave, or a signal modulated by a carrier, over a transmission medium or communication link. The  
20 machine-readable medium may include any medium that can store or transfer information in a form readable and executable by a machine (e.g. a processor, a computer, etc.). Examples of the machine-readable medium include an electronic circuit, a semiconductor memory device, a ROM, a flash memory, an erasable programmable ROM (EPROM), a floppy diskette, a compact disk CD-ROM, an optical  
25 disk, a hard disk, a fiber optic medium, a radio frequency (RF) link, etc. The computer data signal may include any signal that can propagate over a transmission medium such as electronic network channels, optical fibers, air, electromagnetic, RF links, etc. The code segments may be downloaded via networks such as the Internet, Intranet, etc.

While embodiments of the invention have been described with reference to illustrative embodiments, these descriptions are not intended to be construed in a limiting sense. Various modifications of the illustrative embodiments, as well as other embodiments of the invention, which are apparent to persons skilled in the art to which

5   embodiments of the invention pertain, are deemed to lie within the spirit and scope of the invention.